

From the Rendering Equation to Stratified Light Transport Inversion



*Tian-Tsong Ng
Ramanpreet Singh Pahwa
Jiamin Bai
Kar-Han Tan
Ravi Ramamoorthi*

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2010-151
<http://www.eecs.berkeley.edu/Pubs/TechRpts/2010/EECS-2010-151.html>

December 9, 2010

Report Documentation Page		<i>Form Approved OMB No. 0704-0188</i>
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.		
1. REPORT DATE 24 NOV 2010	2. REPORT TYPE	3. DATES COVERED 00-00-2010 to 00-00-2010
4. TITLE AND SUBTITLE From the Rendering Equation to Stratified Light Transport Inversion		5a. CONTRACT NUMBER
		5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER
		5e. TASK NUMBER
		5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of California, Berkely, Electrical Engineering and Computer Sciences, Berkeley, CA, 94720		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited		
13. SUPPLEMENTARY NOTES		
14. ABSTRACT Recent advances in fast light transport acquisition have motivated new applications for forward and inverse light transport. While forward light transport enables image relighting, inverse light transport provides new possibilities for analyzing and cancelling interreflections, to enable applications like projector radiometric compensation and light bounce separation. With known scene geometry and diffuse reflectance, inverse light transport can be easily derived in closed form. However with unknown scene geometry and reflectance properties, we must acquire and invert the scene's light transport matrix to undo the effects of global illumination. For many photometric setups such as that of a projector-camera system, the light transport matrix often has a size of 105 105 or larger. Conventional inversion techniques, such as the pseudo-inverse, are accurate but impractical computationally at these resolutions. In this work, we explore a theoretical analysis of inverse light transport, relating it to its forward counterpart, expressed in the form of the rendering equation. It is well known that forward light transport has a Neumann series that corresponds to adding bounces of light. In this paper, we show the existence of a similar inverse series, that zeroes out the corresponding physical bounces of light. We refer to this series solution as stratified light transport inversion, since truncating to a certain number of terms corresponds to cancelling the corresponding interreflection bounces. The framework of stratified inversion is general and may provide insight for other problems in light transport and beyond, that involve large-size matrix inversion. It is also efficient, requiring only sparse matrix-matrix multiplications. Our practical application is to radiometric compensation, where we seek to project patterns onto real-world surfaces, undoing the effects of global illumination. We use stratified light transport inversion to efficiently invert the acquired light transport matrix for a static scene, after which interreflection cancellation is a simple matrix-vector multiplication to compensate the input image for projection.		
15. SUBJECT TERMS		

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 31	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

Copyright © 2010, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

We are grateful to Tony Quee-Seng Quek for giving feedbacks on an earlier draft of the manuscript. We wish to thank Manmohan Chandraker and Mao-Pei Tsui for their great suggestions on various technical issues. We would also like to thank Yasuyuki Matsushita, Jiaping Wang and Yue Dong of Microsoft Research Asia for their discussion and suggestions on experimental techniques. Finally, we would like to thank Zhiyong Huang and Joo Hwee Lim for their support on this project. Ramanpreet Singh Pahwa and Jiamin Bai are supported by National Science Scholarships awarded by A*STAR Graduate Academy of Singapore. Ravi Ramamoorthi is supported by ONR YIP grant N00014-10-1-0032, as well as NSF CAREER Award 09-24968 and ONR PECASE grant N00014-09-1-0741.

From the Rendering Equation to Stratified Light Transport Inversion

Tian-Tsong Ng Ramanpreet Singh Pahwa Jiamin Bai
Kar-Han Tan Ravi Ramamoorthi

November 24, 2010

Abstract

Recent advances in fast light transport acquisition have motivated new applications for forward and inverse light transport. While forward light transport enables image relighting, inverse light transport provides new possibilities for analyzing and cancelling interreflections, to enable applications like projector radiometric compensation and light bounce separation. With known scene geometry and diffuse reflectance, inverse light transport can be easily derived in closed form. However, with unknown scene geometry and reflectance properties, we must acquire and invert the scene’s light transport matrix to undo the effects of global illumination. For many photometric setups such as that of a projector-camera system, the light transport matrix often has a size of $10^5 \times 10^5$ or larger. Conventional inversion techniques, such as the pseudo-inverse, are accurate but impractical computationally at these resolutions.

In this work, we explore a theoretical analysis of inverse light transport, relating it to its forward counterpart, expressed in the form of the rendering equation. It is well known that forward light transport has a Neumann series that corresponds to adding bounces of light. In this paper, we show the existence of a similar inverse series, that zeroes out the corresponding physical bounces of light. We refer to this series solution as stratified light transport inversion, since truncating to a certain number of terms corresponds to cancelling the corresponding interreflection bounces. The framework of stratified inversion is general and may provide insight for other problems in light transport and beyond, that involve large-size matrix inversion. It is also efficient, requiring only sparse matrix-matrix multiplications. Our practical application is to radiometric compensation, where we seek to project patterns onto real-world surfaces, undoing the effects of global illumination. We use stratified light transport inversion to efficiently invert the acquired light transport matrix for a static scene, after which interreflection cancellation is a simple matrix-vector multiplication to compensate the input image for projection.

1 Introduction

Global illumination and interreflection effects are important aspects of real-world scenes. For forward light transport simulation, a theoretical foundation in terms of the rendering equation by Kajiya (1986) is now well established. More recently, it has become




Figure 1: An illustration of radiometric compensation with stratified inverses on a complex scene with an array of hemispheres.

possible to directly acquire the light transport for a real scene, that in essence physically computes the rendering equation for different lighting conditions. Forward light transport can be viewed as a linear operator that maps a lighting configuration to an image sensor. For instance, in a projector-camera system, the light transport maps a projector input to a camera output. Recent approaches have enabled great efficiency and generality in light transport acquisition as shown in Debevec et al (2000), Sen et al (2005), Ding et al (2009), Peers et al (2009) and Wang et al (2009). Even for purely synthetic scenes, precomputed light transport is now popular for rendering as shown in Ramamoorthi (2007).

Acquired or precomputed light transport matrices have so far been used primarily for relighting, that mathematically now becomes a simple matrix-vector multiplication. However, there is also a large class of applications enabled by *inverse light transport*. One possibility, studied by Seitz et al (2005), is stratifying the image into the different bounces of light (direct lighting, first interreflection bounce and so on). In this paper, we use projector radiometric compensation as the motivating application (Habe et al, 2007; Wetzstein and Bimber, 2007; Ding et al, 2009; Ng et al, 2009). Our goal is to project images and patterns from the projector onto uncalibrated real-world scenes with unknown geometry and complex reflectance properties. In these cases, the observed image includes interreflection effects, that we seek to compensate to obtain the desired result. If we could invert the light transport matrix, we can compensate the input, so that the observed image upon projection matches the desired image. However, even the lowest resolutions of real-world cameras and projectors usually cause the light transport matrix to be of size at least $10^5 \times 10^5$ —that makes direct matrix inversion, such as computing the pseudo-inverse, computationally intractable.

While forward light transport has been well studied, little is known about the theoretical and computational structure and properties of inverse light transport. Furthermore, while it is widely known that forward light transport can be incrementally computed by adding light bounces through a Neumann series, it remains unknown if such an incremental computation has a counterpart for inverse light transport. In this paper, we address these fundamental questions, identifying a structure of the inverse light transport derived directly from the rendering equation, which mirrors that of its forward counterpart. We show that inverse light transport can be expressed in a polynomial series similar to the forward Neumann series, and that each term or iteration zeroes out the corresponding physical light bounce. Such a series expression presents a natural structure for the inverse light transport to be computed iteratively, which gives rise to a

stratified inversion algorithm.

In practical terms, stratified inverses enable projector radiometric compensation at moderate to high scene resolutions, once the light transport of a static scene has been acquired and without assuming known scene geometry. They provide a natural framework for gradual approximation of inverse light transport, with only a small number of sparse matrix-matrix multiplications required. We demonstrate results (see Fig. 12) with image resolution of 384×512 , resulting in a light transport matrix of size 196608×196608 . The 4-bounce stratified inverse can be computed in under a minute, while memory and computational issues make even running a direct pseudoinverse impossible on our hardware at this resolution. We emphasize that the full light transport matrix is inverted, that then allows a simple matrix-vector multiplication to compensate each image that is projected, just like in (forward) image relighting. Bimber et al (2006) and Bimber (2006) explored a similar approximation framework for light transport inversion but under the assumption of known scene geometry.

This paper is an extension of Ng et al (2009) in the following aspects. First, we propose an insightful way to derive the stratified inverse light transport directly from Kajiya’s rendering equation. This derivation makes explicit the connection between the iterative computation framework of Bimber et al (2006) and Bimber (2006) with known scene geometry and the stratified inversion framework of Ng et al (2009) with only light transport measurement as input. Second, we show a Neumann interpretation for the stratified inverses in terms of physical bounces of light, which brings out an interesting correspondence between the forward and the inverse light transport (a brief summary of this result is described by us in Bai et al (2010), but this paper presents the full derivation and analysis.) Although Seitz et al (2005) has showed that inverse light transport can be used for separating light bounces in forward light transport, the physical meaning of the polynomial terms in inverse light transport is novel. Third, we show that, by interpreting stratified inversion in a numerical preconditioning framework, the exact knowledge of the first-bounce light transport (which is in general unknown) is not required for the convergence of stratified inversion. Hence, for diffuse scenes, the first bounce light transport can be computed using a result in Seitz et al (2005), by applying stratified inverse on the potentially large-size light transport. With accurate estimation of the first bounce light transport, stratified inverses correspond to physical light bounces. We showed the accurate estimation of a first bounce light transport in a simulation in Subsection 5.1.

The paper is organized as follows. We start with a review of previous work in Section 2. In Section 3, we go through the basics of the rendering equation and forward light transport. The main contribution of this work is described in Section 4 where the stratified inverse is derived from the rendering equation and given a Neumann interpretation of bounce cancellation. In Section 5, we validate the computational framework through a simple example. In Section 6, stratified inverses are experimented with for a projector radiometric compensation application on real scenes. Finally, we conclude and discuss future work directions in Section 7.

Table 1: Comparison of Prior Works for Inverse Light Transport

	Analytic Formula	Iterative Computation	Matrix Simplification	Matrix Stratification	Theoretical Analysis
Known Scene Geometry	Mukaigawa et al (2006)	Bimber et al (2006); Bimber (2006)	-	-	-
Known Scene Geometry	-	Bai et al (2010)	Habe et al (2007); Wetzstein and Bimber (2007); Ding et al (2009)	This work; Ng et al (2009)	This work; Seitz et al (2005); Bai et al (2010)

2 Previous Work

We first discuss radiometric compensation, where most of the previous work on computational methods for inverse light transport has been developed. This will form our motivating practical application. We then briefly review relevant literature in light transport acquisition, global illumination rendering for forward light transport, and inverse light transport.

2.1 Radiometric Compensation

Radiometric compensation has a long history, with both geometric and photometric distortions considered. In this paper, we assume geometric calibration and focus on photometric issues.

One to One Mapping with No Interreflections: Most early work considered the mapping between projector and camera to be one-to-one (Raskar et al, 1998; Majumder et al, 1999; Nayar et al, 2003; Bimber et al, 2005; Fujii et al, 2005; Ashdown et al, 2006; Song and Cham, 2005). For a comprehensive survey, please refer to Yang et al (2004) and Raskar et al (1999). Although the inversion of such a mapping is trivial, the one-to-one mapping can neither capture the global illumination effects such as interreflection, refraction and scattering in the scene, nor compensate for them.

Iterative Methods for Each Input: For the case of known scene geometry (and Lambertian reflectance), Bimber et al (2006) and Bimber (2006) showed that the light interreflection of a concave scene can be compensated by computing the correct input to the projector iteratively. This is done using Jacobi iteration on the projector input vector, with the form factor matrix derived from the scene geometry. In the case of unknown scene geometry, but given the light transport of the scene, we showed in Bai et al (2010) that projector radiometric compensation can be similarly computed

iteratively. These approaches relate closely to the radiosity method for diffuse global illumination in forward rendering (Hanrahan et al, 1991; Gortler et al, 1993).

These iterative methods are effectively solving a linear equation without the need for direct matrix inversion—hence they are more efficient when only a single image needs to be compensated. However, for multiple images or the case of displaying a video from a projector on a static scene, the compensated projector input needs to be separately and iteratively computed for every frame. This can require considerable computation, which may preclude real-time frame rates. While stratified inverses are also conceptually similar as an iterative approach, our method (pre)computes the full matrix inverse only once. Thereafter, the compensated projector input can be obtained with a single-step matrix-vector multiplication for any desired image. In effect, we tradeoff more initial computation of a full matrix inverse (that is still however very efficient, taking less than a minute) for a very simple non-iterative compensation method at run-time. Our formulation also derives directly from the rendering equation, and the theory encompasses general non-Lambertian materials.

Analytic Solution for Known Geometry: In fact, with a known Lambertian form factor, Mukaigawa et al (2006) showed that the compensated projector input can be easily computed in closed form without iterative computation. A similar conceptual observation was made by Seitz et al (2005). However, the assumption of known form factor is not realistic except for the scenes derived from a graphics model or ones with simple geometry and reflectance which can easily be measured. Otherwise, it is more realistic to measure the light transport of a scene directly for a projector-camera system.

Approximations to Full Inverse Light Transport: Projector radiometric compensation can be achieved through the inverse light transport, which can only be computed by inverting the light transport matrix \mathbf{T} as in Habe et al (2007), Wetzstein and Bimber (2007) and Ding et al (2009). The real challenge for the matrix inversion is the enormous size of a typical \mathbf{T} , given the resolutions of projector and camera. In order to produce \mathbf{T} of a size feasible for pseudo-matrix inversion, Habe et al (2007) downsized the full-size \mathbf{T} through grouping the neighboring projector pixels into single super-pixels. However, the method only works well for a scene with relatively simple geometry and texture of lower spatial frequency.

Another approach is to sacrifice the fidelity of \mathbf{T} by simplifying the matrix without changing its size. Ding et al (2009) introduced a constraint known as the *display constraint*, such that a camera pixel can only receive light from a single projector pixel. For \mathbf{T} to comply with the display constraint, each row of \mathbf{T} should only retain the largest element by setting the other elements to zero. In this manner, the display constraint will make the resultant \mathbf{T} column-orthogonal and matrix inversion can be obtained through column-wise computation. However, the interreflection effect in \mathbf{T} is effectively removed, and this creates a discrepancy between the actual light transport and the display constraint compliant \mathbf{T} . As a result, projector compensation using such a matrix cannot remove any global illumination effect.

Wetzstein and Bimber (2007) simplified \mathbf{T} by clustering links between projector and camera pixels to form clusters which correspond to independent sub-matrices. Matrix inversion can be computed separately for each cluster and is efficient if each

individual cluster is of manageable size. Naturally, independent clusters may not exist and some of the links between projector and camera pixels often have to be severed in order to produce clusters of manageable size. This spatial clustering method effectively ignores interreflections between clusters and reduces effectiveness of projector compensation.

Comparison and Summary: Table 1 summarizes the prior work in light-transport-based projector radiometric compensation, including our own previous conference papers (Ng et al, 2009; Bai et al, 2010). For known scene geometry and known Lambertian reflectance, simple analytic approaches are available as shown in Mukaigawa et al (2006). However, for unknown scene geometry, such methods are not possible. Fast iterative computation methods can be used in Bimber et al (2006), Bimber (2006) and Bai et al (2010), but must be rerun for each input image. Moreover, many previous methods assume Lambertian reflectance as in Mukaigawa et al (2006), Bimber et al (2006) and Bimber (2006), while our formulation considers general materials and derives directly from the rendering equation. Alternatively, simplifications can be made to enable full matrix inversion as in Habe et al (2007), Ding et al (2009), Wetzstein and Bimber (2007). However, these approaches introduce approximations, and no principled analysis of error or convergence properties is available.

The contributions of this paper are primarily theoretical, with an important practical advance for radiometric compensation. Most importantly, we develop a principled theoretical analysis of inverse light transport. We build on the seminal work of Seitz et al (2005) but go much further in exploring the relationships to forward rendering and the rendering equation. In particular, we develop a stratified inverse method, that cancels interreflection bounces in analogy to the forward Neumann series that adds bounces of light. In practical terms, this provides a principled and efficient way to compute the full matrix inverse in a stratified fashion, using only sparse matrix-matrix multiplications. Once the inverse of the light transport matrix is computed, compensation is direct, with a simple matrix-vector multiplication (of the inverse matrix and desired image).

2.2 Light Transport Acquisition

A large body of work over the last decade in computer graphics and computer vision has dealt with acquiring light transport matrices, indicating how a real scene responds to light from all directions or projector pixels (Debevec et al, 2000; Sen et al, 2005; Ding et al, 2009; Peers et al, 2009; Wang et al, 2009). While in this paper we focus on real scenes, the use of precomputed transport is increasingly common even for synthetic scenes, in applications like real-time relighting (Ng et al, 2003).

There are many existing methods to acquire transport of a scene. The brute-force method turns on the projector pixels one by one while the response of each projector pixel is captured by a camera, and forms a column in \mathbf{T} . This method produces an accurate \mathbf{T} but it is time consuming and storage intensive as the number of images that need to be acquired is equivalent to the number of projector pixels. In Sen et al (2005), a multi-resolution and adaptive method was proposed to measure the transport. In Peers et al (2009) and Sen and Darabi (2009), a compressive sensing approach was proposed to exploit sparsity in \mathbf{T} , and compute the response of each pixel by projecting

patterned illumination. A fast method in Wang et al (2009) was introduced recently to take advantage of the coherency in the rows and columns of \mathbf{T} , given a specific hardware setup to capture images from the projector viewpoints. In Ding et al (2009), a deterministic stripe-scanning method was proposed to acquire \mathbf{T} where horizontal and vertical stripes scan through the scene in a sequential manner. This method is efficient and simple, but it tends to consistently over-estimate the projector pixel response. For a comprehensive review of light transport acquisition, we refer readers to Peers et al (2009), Wang et al (2009).

In this paper, we simply use these methods to acquire the original light transport, as our main focus is on light transport *inversion*.

2.3 Global Illumination Rendering

Much of our inspiration draws from the thorough study in computer graphics of the forward problem, or global illumination rendering. In particular, our stratified inverse light transport is derived directly from Kajiya’s rendering equation (Kajiya, 1986) using the operator notation in Arvo et al (1994). While we currently simply use sparse matrices to represent \mathbf{T} , we are also interested in exploring connections with hierarchical and wavelet radiosity as in Hanrahan et al (1991) and Gortler et al (1993) in future; our stratified matrix inversion also bears some conceptual similarities to Jacobi and Gauss-Seidel iterative methods used in radiosity.

2.4 Inverse Light Transport

Previous work on inverse rendering has considered inversion of the direct reflection equation to acquire lighting and reflectance properties as in (Marschner, 1998; Ramamoorthi and Hanrahan, 2001). Yu et al (1999) developed an inverse global illumination method for BRDF estimation. However, all these methods assume the scene geometry is known, and usually work with lower resolutions for lighting, which makes analysis of interreflections much easier (and often requires only a few input images). In contrast, our setup is closer to Seitz et al (2005), where only the light transport matrix is observed—both geometry and reflectance are unknown, and are not explicitly estimated.

Given an inverse light transport matrix, the input illumination that produces a given photo of a scene can be computed. O’Toole and Kutulakos (2010) showed that the input illumination can be computed optically without explicitly measuring the light transport. By simulating matrix-vector multiplication optically, the optical computation is implicitly solving a matrix inversion problem through iterative optical measurements. With K rounds of optical computation, the method computes the input illumination from the K -rank approximation of the light transport. Therefore, this method is efficient for a scene with a lower rank light transport matrix. As it requires four optical measurements for each round of optical computation, it may be inefficient for a scene with a maximal rank light transport matrix. In this work, we are interested in computing an inverse light transport matrix explicitly and focus on maximal-rank light transport which is commonly encountered for projector radiometric compensation.

While not focused on inverse light transport per se, the fast separation of direct and global components by Nayar et al (2006) is also relevant. In theory, by just projecting a high-frequency checker-board pattern and its complement onto the scene, the direct reflection image can be extracted. In practice, due to the resolution limits imposed by the projector and the camera, 25 images were used in Nayar et al (2006). This method has been used primarily for separating a single image (rather than the full transport matrix) into its components, and is suitable when \mathbf{T} is unknown. The method also assumes that the global illumination component is low-frequency, which can be violated in cases of strong localized subsurface or reflection effects.

3 Forward Light Transport

We now introduce the basic formulation of the rendering equation and operator solution for forward light transport and global illumination. Then, in Section 4, we proceed to develop our theoretical formulation of stratified inverse light transport.

The rendering equation (Kajiya, 1986) can be written,

$$L_{out}(\mathbf{x}, \omega_0) = L_e(\mathbf{x}, \omega_0) + \int_M \rho(\mathbf{x}, \omega_i, \omega_0) L_{out}(\mathbf{y}, -\omega_i) V(\mathbf{x}, \mathbf{y}) \frac{\cos \theta_i \cos \theta_o}{\|\mathbf{x} - \mathbf{y}\|^2} dA_y, \quad (1)$$

where $L_{out}(\mathbf{x}, \omega_0)$ is the reflected or outgoing radiance, L_e is the emission corresponding to light sources, ρ is the Bi-directional Reflectance Distribution Function (BRDF) of the scene surface, and V is the binary visibility function. The variables \mathbf{x} , ω_i , ω_0 are respectively the spatial location, incident and outgoing angles on a surface. The visibility function $V(\mathbf{x}, \mathbf{y})$ is 1 if \mathbf{x} and \mathbf{y} are connected by a line of sight, 0 otherwise. The integral is over the area M of all scene surfaces, and weighted by a purely geometric factor involving cosines of incident and outgoing angles, and the distance between \mathbf{x} and \mathbf{y} .

3.1 Operator Solution of Rendering Equation

Following Arvo et al (1994), the rendering equation can be written in operator notation (or equivalently in a discrete matrix form) as,

$$\mathbf{I}_{out} = \mathbf{I}_d + \mathbf{KG}\mathbf{I}_{out}, \quad (2)$$

where \mathbf{I}_{out} is a vector of $L_{out}(\mathbf{x}, \omega_0)$, \mathbf{I}_d is a vector of $L_e(\mathbf{x}, \omega_0)$, \mathbf{G} is a purely geometric operator that takes outgoing or reflected radiance and propagates it within the scene to obtain incident radiance, and \mathbf{K} is a local linear reflection operator that takes incident radiance and turns it into reflected light based on the BRDF of the surface ρ . Note that in operator form as in (2), the formulation holds for general materials, and is not limited to Lambertian surfaces or the radiosity formulation. From (2), we can obtain

$$\mathbf{I}_{out} = (\mathbf{I} - \mathbf{A})^{-1} \mathbf{I}_d, \text{ where } \mathbf{A} = \mathbf{KG}. \quad (3)$$

3.2 Adaptation for Projector-Camera Systems

As opposed to the conventional global illumination formulation, we do not have emissive surfaces per se, but emission is induced by projection onto the scene. Assuming the camera does not see the projector directly, we can replace \mathbf{l}_d with the effective emission, that corresponds to the direct reflection from the projector,

$$\mathbf{l}_d = \mathbf{F}\mathbf{l}_{in}, \quad (4)$$

where \mathbf{l}_{in} is the incident lighting from the projector, and \mathbf{F} is the light transport matrix that corresponds to the first-bounce reflection. Hence, we have,

$$\mathbf{l}_{out} = (\mathbf{I} - \mathbf{A})^{-1}\mathbf{F}\mathbf{l}_{in}. \quad (5)$$

Note that (5) is in the same canonical form as the forward light transport equation $\mathbf{l}_{out} = \mathbf{T}\mathbf{l}_{in}$, with \mathbf{T} now being defined as

$$\mathbf{T} = (\mathbf{I} - \mathbf{A})^{-1}\mathbf{F}. \quad (6)$$

By defining

$$\mathbf{S} = (\mathbf{I} - \mathbf{A})^{-1}, \quad (7)$$

we can write $\mathbf{T} = \mathbf{SF}$ and $\mathbf{l}_{out} = \mathbf{Sl}_d$. It is well known that the expression $\mathbf{S} = (\mathbf{I} - \mathbf{A})^{-1}$ can be expanded in a Neumann series as corresponding intuitively to increasing numbers of bounces of light or interreflections,

$$\mathbf{S} = \mathbf{I} + \mathbf{A} + \mathbf{A}^2 + \mathbf{A}^3 + \dots \quad (8)$$

4 Theory

In this section, we will present the main results of our paper. We first derive the Neumann series for inverse light transport, analogous to (8). We then derive the stratified inverse method, and a physical interpretation of the Neumann series as cancelling the corresponding interreflection bounces of light. We also relate our analysis to numerical techniques based on preconditioning.

4.1 Neumann Series for Inverse Light Transport

Our goal is now to derive an expression for the inverse light transport matrix, \mathbf{T}^{-1} . To do so, it will first be convenient to define another linear operator or matrix \mathbf{R} that accounts only for global illumination or the global component \mathbf{l}_g ,

$$\mathbf{l}_{out} = \mathbf{l}_d + \mathbf{l}_g = \mathbf{l}_d + \mathbf{R}\mathbf{l}_d = (\mathbf{I} + \mathbf{R})\mathbf{l}_d. \quad (9)$$

From $\mathbf{l}_{out} = \mathbf{Sl}_d$, we can write $\mathbf{S} = \mathbf{I} + \mathbf{R}$ and expand \mathbf{S}^{-1} in a Neumann series as

$$\mathbf{S}^{-1} = (\mathbf{I} + \mathbf{R})^{-1} = \mathbf{I} - \mathbf{R} + \mathbf{R}^2 - \mathbf{R}^3 + \dots \quad (10)$$

We can now rewrite (10) simply in terms of \mathbf{T} and \mathbf{F} by substituting $\mathbf{R} = \mathbf{S} - \mathbf{I} = \mathbf{TF}^{-1} - \mathbf{I}$

$$\mathbf{S}^{-1} = \sum_{k=0}^{\infty} (-\mathbf{R})^k = \sum_{k=0}^{\infty} (\mathbf{I} - \mathbf{TF}^{-1})^k, \quad (11)$$

where $(-\mathbf{R})^0 = \mathbf{I}$.

Finally, from $\mathbf{T} = \mathbf{SF}$, we can write the inverse light transport \mathbf{T}^{-1} as

$$\mathbf{T}^{-1} = \mathbf{F}^{-1} \mathbf{S}^{-1} = \mathbf{F}^{-1} \sum_{k=0}^{\infty} (\mathbf{I} - \mathbf{TF}^{-1})^k, \quad (12)$$

which will converge if $\|\mathbf{I} - \mathbf{TF}^{-1}\| < 1$.

4.2 Stratified Inverses

From (12), we can define a series of approximations to \mathbf{T}^{-1} by dropping the higher-order terms one by one. We call these approximations the *stratified inverse light transport* of the scene and denote them as \mathbf{T}_n^{-1} ,

$$\mathbf{T}_n^{-1} \doteq \mathbf{F}^{-1} \sum_{k=0}^n (\mathbf{I} - \mathbf{TF}^{-1})^k. \quad (13)$$

We now derive further insights using the binomial theorem to expand in terms of \mathbf{TF}^{-1} . In particular,

$$\begin{aligned} \mathbf{T}_n^{-1} &= \mathbf{F}^{-1} \sum_{l=0}^n \left(\sum_{k=l}^n \binom{k}{l} (-\mathbf{TF}^{-1})^l \right) \\ &= \mathbf{F}^{-1} \sum_{l=0}^n \binom{n+1}{l+1} (-\mathbf{TF}^{-1})^l, \end{aligned} \quad (14)$$

where the last line uses a well known combinatorial summation identity. This clearly shows that the stratified inverse is a polynomial in terms of \mathbf{TF}^{-1} .

4.3 Physical Interpretation: Inverse Neumann Series as Cancelling Physical Bounces of Light

So far, we have introduced the stratified inverse, that is the inverse analog to the forward Neumann series. However, while the forward Neumann series corresponds physically to adding bounces of light, it is not clear what physical interpretation the inverse Neumann series has. We will now derive a perhaps surprising result—just as each term in the forward series adds a physical bounce of light, each term in the inverse series cancels the corresponding bounce. However, convergence is oscillatory in the inverse series, owing to the alternating negative and positive signs in (10). Because of this, coefficients of higher bounces will oscillate until they are zeroed by the corresponding term in the inverse series.

We start with the basic relations of (7) and (10), that $\mathbf{S} = (\mathbf{I} - \mathbf{A})^{-1}$ and $\mathbf{S}^{-1} = (\mathbf{I} + \mathbf{R})^{-1}$,

$$\mathbf{R} = \mathbf{A} + \mathbf{A}^2 + \dots = \mathbf{A}(\mathbf{I} - \mathbf{A})^{-1}. \quad (15)$$

Now, from (15) above, we can derive a series,

$$\mathbf{S}^{-1} = \sum_{k=0}^{\infty} (-1)^k \mathbf{R}^k = \sum_{k=0}^{\infty} (-1)^k [\mathbf{A}(\mathbf{I} - \mathbf{A})^{-1}]^k. \quad (16)$$

While in general, raising a matrix (or operator) product to a power is complicated because of non-commutativity, in our case everything involves powers of \mathbf{A} , and so \mathbf{A} and $(\mathbf{I} - \mathbf{A})^{-1}$ commute, and can be exponentiated separately. We can now put this together to derive,

$$\mathbf{S}_n^{-1} = \sum_{k=0}^n (-1)^k \mathbf{A}^k (\mathbf{I} - \mathbf{A})^{-k}. \quad (17)$$

4.3.1 Binomial Series Expansion

Using a standard binomial series expansion for $(\mathbf{I} - \mathbf{A})^{-k}$, this can be written as

$$\mathbf{S}_n^{-1} = \sum_{k=0}^n (-1)^k \mathbf{A}^k \sum_{l=0}^{\infty} \binom{k+l-1}{l} \mathbf{A}^l. \quad (18)$$

Our next step is to combine the powers of \mathbf{A} , using that $m = l + k$ and $l = m - k$,

$$\mathbf{S}_n^{-1} = \sum_{k=0}^n \sum_{m=k}^{\infty} (-1)^k \binom{m-1}{m-k} \mathbf{A}^m. \quad (19)$$

It will simplify the later analysis if we treat $k = 0$ as a special case, given obviously from (17) as the identity. We also use $(m-1) - (m-k) = (k-1)$ in the combination,

$$\mathbf{S}_n^{-1} = \mathbf{I} + \sum_{k=1}^n \sum_{m=k}^{\infty} (-1)^k \binom{m-1}{k-1} \mathbf{A}^m. \quad (20)$$

To proceed further, we need to transpose the order of the summations. The outer summation should be about m , which controls the powers. It is clear that we require $m \geq k$, which in turn leads to the relations $k \leq m$ and (because we are considering the n term inverse series) that $k \leq n$,

$$\mathbf{S}_n^{-1} = \mathbf{I} + \sum_{m=1}^{\infty} \left[\sum_{k=1}^{\min(m,n)} (-1)^k \binom{m-1}{k-1} \right] \mathbf{A}^m. \quad (21)$$

4.3.2 Base Cases

We treat the simple cases when $n = 0, 1$ and $m = 1$ first. When $n = 0$, the expression above just reduces to the identity (no bounce is cancelled as expected). When $n = 1$, only the $k = 1$ term is relevant, so we have

$$\mathbf{S}_1^{-1} = \mathbf{I} - \sum_{m=1}^{\infty} \mathbf{A}^m, \quad (22)$$

Table 2: Coefficients of \mathbf{S}_n^{-1} and $\mathbf{S}_n^{-1}\mathbf{S}$

\mathbf{S}_n^{-1}	I	A	A²	A³	A⁴	A⁵	A⁶	A⁷
\mathbf{S}_0^{-1}	1	0	0	0	0	0	0	0
\mathbf{S}_1^{-1}	1	-1	-1	-1	-1	-1	-1	-1
\mathbf{S}_2^{-1}	1	-1	0	1	2	3	4	5
\mathbf{S}_3^{-1}	1	-1	0	0	-1	-3	-6	-10
\mathbf{S}_4^{-1}	1	-1	0	0	0	1	4	10
\mathbf{S}_5^{-1}	1	-1	0	0	0	0	-1	-5
\mathbf{S}_6^{-1}	1	-1	0	0	0	0	0	1
\mathbf{S}_7^{-1}	1	-1	0	0	0	0	0	0

$\mathbf{S}_n^{-1}\mathbf{S}$	I	A	A²	A³	A⁴	A⁵	A⁶	A⁷
$\mathbf{S}_0^{-1}\mathbf{S}$	1	1	1	1	1	1	1	1
$\mathbf{S}_1^{-1}\mathbf{S}$	1	0	-1	-2	-3	-4	-5	-6
$\mathbf{S}_2^{-1}\mathbf{S}$	1	0	0	1	3	6	10	15
$\mathbf{S}_3^{-1}\mathbf{S}$	1	0	0	0	-1	-4	-10	-20
$\mathbf{S}_4^{-1}\mathbf{S}$	1	0	0	0	0	1	5	15
$\mathbf{S}_5^{-1}\mathbf{S}$	1	0	0	0	0	0	-1	-6
$\mathbf{S}_6^{-1}\mathbf{S}$	1	0	0	0	0	0	0	1
$\mathbf{S}_7^{-1}\mathbf{S}$	1	0	0	0	0	0	0	0

Note: The series, \mathbf{S}_n^{-1} and $\mathbf{S}_n^{-1}\mathbf{S}$, exhibit oscillatory convergence towards $\mathbf{I} - \mathbf{A}$ and \mathbf{I} respectively. The n term series is accurate up to \mathbf{A}^n , and in fact cancels or zeroes bounces up to that order, with errors only in higher-order terms or bounces $n + 1$ and higher.




Figure 2: Coefficients of \mathbf{S}_n^{-1} for $m = 0, \dots, 7$.

where we note that for $k = 1$, the $k - 1$ term in the combination reduces it to 1, and $(-1)^k = -1$. This is indeed the expected result, since $\mathbf{S}_1^{-1} = \mathbf{I} - \mathbf{R}$, and we know from (15) that $\mathbf{R} = \mathbf{A} + \mathbf{A}^2 + \dots$

Finally, the special case $m = 1$ will be useful. In this case (assuming $n > 1$), the second summation in (21) will have upper limit $m = 1$, and the coefficient will simply be 1. Thus, for $n > 1$ (the cases $n = 0$ and $n = 1$ have already been dealt with),

$$\mathbf{S}_n^{-1} = \mathbf{I} - \mathbf{A} + \sum_{m=2}^{\infty} \left[\sum_{k=1}^{\min(m,n)} (-1)^k \binom{m-1}{k-1} \right] \mathbf{A}^m. \quad (23)$$

4.3.3 Zeroing of Higher-Order Bounces

Now, consider the case when $m \leq n$. In this case, the second summation has a limit of $m > 1$, and the coefficient of \mathbf{A}^m becomes

$$\sum_{k=1}^m (-1)^k \binom{m-1}{k-1} = - \sum_{k'=0}^{m'} (-1)^{k'} \binom{m'}{k'} = 0, \quad (24)$$

where $m' = m - 1$ and $k' = k - 1$ (note this only works for $m > 1$; the $m = 1$ term is given as a special case in (23)). The expression above is clearly 0, since those are the coefficients in a binomial expansion of the expression $(1+x)^{m'}$, with $x = -1$.

This implies a key result, that the \mathbf{A}^m terms vanish for $2 \leq m \leq n$, which in turn implies that

$$\mathbf{S}_n^{-1} = \mathbf{I} - \mathbf{A} + O(\mathbf{A}^{n+1}), \quad (25)$$

where $O(\cdot)$ denotes higher order terms, and $n > 1$. Note that since $\mathbf{S} = (\mathbf{I} - \mathbf{A})^{-1}$, the final result we desire¹ is simply $\mathbf{S}^{-1} = \mathbf{I} - \mathbf{A}$. Equation (25) states that terms up to order n are correct, and in fact terms from $[\mathbf{A}^2 \dots \mathbf{A}^n]$ are 0. Note however that the terms (bounces) \mathbf{A}^{n+1} and higher oscillate and are not zeroed until the corresponding inverse series term is considered.

¹In fact, the analytic solution of Mukaigawa et al (2006) essentially uses this observation in cases where the form factors and hence \mathbf{A} are known. In our case, \mathbf{A} is unknown, and recovering it is essentially equivalent to inverting the light transport.




Figure 3: Coefficients of $S_n^{-1}S$ for $m = 0, \dots, 7$.

4.3.4 Bounce Cancellation

We have seen how higher-order terms are zeroed in the inverse operator series. We now show that applying the n -term inverse series to the original cancels the first n bounces. For this we write,

$$S_n^{-1}S = [(I - A) + O(A^{n+1})][I - A]^{-1}. \quad (26)$$

It is clear that the first part $I - A$ creates the identity as desired. The product $O(A^{n+1})(I - A)^{-1}$ is still of $O(A^{n+1})$, since the inverse can be expanded in a Neumann series. Therefore,

$$S_n^{-1}S = I + O(A^{n+1}). \quad (27)$$

In other words, the n -term inverse series annihilates bounces $[1 \dots n]$, leaving only bounces $n + 1$ and higher.

4.3.5 Analytic Forms

In fact, the inner summation in (23) can be performed symbolically (we did so using Mathematica), to derive

$$S_n^{-1} = I - A + (-1)^n \sum_{m=n+1}^{\infty} \binom{m-2}{n-1} A^m \quad (28)$$

$$S_n^{-1}S = I + (-1)^n \sum_{m=n+1}^{\infty} \binom{m-1}{n} A^m. \quad (29)$$

Table 2 shows coefficients for $(m, n) \leq 7$. Fig. 2 and Fig. 3 show the coefficients of S_n^{-1} and $S_n^{-1}S$ respectively for $n \leq 12$ and $m \leq 7$. Owing to the $(-1)^n$ term, these coefficients oscillate until they are zeroed.

4.4 Interpretation as Numerical Preconditioner

We have so far derived the stratified inverse method from the rendering equation, allowing us to give a physical interpretation as cancelling interreflection bounces. We

now briefly also show how it can be related to a numerical inverse series with preconditioners. This purely numerical interpretation allows us to relax some conditions for practical implementation.

For any matrix \mathbf{T} , if we can find a matrix (preconditioner) \mathbf{P} that is easy to invert, we can write

$$\mathbf{TP}^{-1} = \mathbf{I} - (\mathbf{I} - \mathbf{TP}^{-1}), \quad (30)$$

from which it follows that

$$\mathbf{T}^{-1} = \mathbf{P}^{-1}(\mathbf{I} - (\mathbf{I} - \mathbf{TP}^{-1}))^{-1}, \quad (31)$$

Expanding this in a Neumann series,

$$\mathbf{T}^{-1} = \mathbf{P}^{-1} \sum_{k=0}^{\infty} (\mathbf{I} - \mathbf{TP}^{-1})^k, \quad (32)$$

with the convergence condition $\|\mathbf{I} - \mathbf{TP}^{-1}\| < 1$. Equation (12) can be viewed as a specialized case where, from physical intuition, we take $\mathbf{P} = \mathbf{F}$ or first bounce light transport. Note that the physical bounce cancellation results just explored require this physical perspective, and cannot be derived for general Neumann series.

The numerical perspective offers two important practical generalizations. First, we can use any matrix as a preconditioner, and need not exactly determine the first bounce transport \mathbf{F} . We simply use a diagonal matrix $\mathbf{P} \approx \mathbf{F} = \text{diag}(\mathbf{T})$. Numerically, this is a Jacobi preconditioner. Physically, it closely approximates the first bounce of light transport for an appropriate parameterization that matches corresponding elements of camera and projector (Seitz et al, 2005), since those are the elements that are nonzero from direct lighting. Moreover, an element does not reflect onto itself, so errors will only be second order. While this is not the exact \mathbf{F} , it suffices numerically.

In fact, since the convergence of stratified inverses does not depend on the exact knowledge of the true \mathbf{F} , it enables us to compute the true \mathbf{F} for diffuse scenes. Seitz et al (2005) showed that, for diffuse scenes, the true \mathbf{F} can be computed exactly provided the light transport matrix can be accurately inverted, as

$$\mathbf{F}_{ii} = \frac{1}{(\mathbf{T}^{-1})_{ii}}, \quad (33)$$

where the subscript index in \mathbf{F}_{ii} indicates the diagonal elements of \mathbf{F} . Such computation is demonstrated in a simple simulation in Section 5. For a large-size light transport, stratified inverses help to break the chicken and egg dependency between the availability of the true \mathbf{F} and the feasibility of computing \mathbf{T}^{-1} . In addition, this also indicates that the physical interpretation of stratified inverses as cancelling interreflection bounces as shown in the previous subsection can be exactly computed for diffuse scenes.

Second, our theoretical framework applies to general non-Lambertian materials, as does the rendering equation from which it is derived. However, the operators do have both spatial and directional dependence, that technically requires us to consider the full light field or space of views. In applying the theory to the common practical setup of a single camera-projector pair, a widely used previous approach is to assume Lambertian




Figure 4: *Simulation scenes:* (a) A Lambertian dihedral with two facets on each of the two panels. The two facets on the left panel subdivide the panel horizontally while the two on the right subdivide the panel vertically. The dotted lines overlaid on the figure show the subdivision for the facets. (b) An open Cornell box with 5 faces.

reflectance as in Seitz et al (2005), in which case the operators can be written without directional dependence. However, the numerical algorithm does not need to make this assumption, even for a single camera-projector pair, since the stratified matrix inversion method still works. For the same reasons, our method is robust to moderate amounts of subsurface scattering, even though the rendering equation theory does not apply to volumetric effects. Our examples, such as the scene in Fig. 1, show both moderate glossiness and subsurface effects. We emphasize that this is a practical issue only; the theory applies directly to general non-Lambertian materials.

5 Simulation

In this section, we evaluate the behavior of the stratified inverse in terms of its convergence properties and rate, as well as its computational accuracy and efficiency versus the true inverse.

5.1 Simple Example

To show the oscillatory convergence of the stratified inverse, we perform a simulation on a $4 \times 4 \mathbf{T}$ matrix of a simple diffuse dihedral scene as shown in Fig. 4(a). Each panel of the dihedral is made up of two facets where the two facets on the left panel subdivide the panel horizontally while the two on the right subdivide the panel vertically. Altogether there are four facets in the entire scene. In this example, we simulate the case where the first-bounce light transport for all the facets is equal to a value 1.750. Hence, the actual first-bounce light transport \mathbf{F} in this example is a diagonal matrix with the diagonal entries being 1.750. On this example light transport, we will compute the corresponding stratified inverses and evaluate the impact of our method in approximating \mathbf{F} with the diagonal of \mathbf{T} , as opposed to simply assuming $\mathbf{F} = \mathbf{I}$, i.e., ignoring \mathbf{F} in the computation.

The 4×4 \mathbf{T} matrix from our four-facet scene with entries rounded to four decimal places is given by:

$$\mathbf{T} = \begin{pmatrix} 1.7713 & 0.0213 & 0.1835 & 0.0675 \\ 0.0213 & 1.7713 & 0.1835 & 0.0675 \\ 0.1835 & 0.1835 & 1.7875 & 0.0137 \\ 0.0675 & 0.0675 & 0.0137 & 1.7550 \end{pmatrix}. \quad (34)$$

In this case, we can easily compute its true inverse and treat it as the ground truth:

$$\mathbf{T}^{-1} = \begin{pmatrix} 0.5715 & 0.0000 & -0.0587 & -0.0220 \\ 0.0000 & 0.5715 & -0.0587 & -0.0220 \\ -0.0587 & -0.0587 & 0.5715 & 0.0045 \\ -0.0220 & -0.0220 & 0.0045 & 0.5715 \end{pmatrix}. \quad (35)$$

As \mathbf{F} is in general unknown given a \mathbf{T} matrix, we approximate \mathbf{F} with the diagonal of \mathbf{T} :

$$\hat{\mathbf{F}} = \begin{pmatrix} 1.7713 & 0.0000 & 0.0000 & 0.0000 \\ 0.0000 & 1.7713 & 0.0000 & 0.0000 \\ 0.0000 & 0.0000 & 1.7875 & 0.0000 \\ 0.0000 & 0.0000 & 0.0000 & 1.7550 \end{pmatrix}. \quad (36)$$

Note that the estimated $\hat{\mathbf{F}}$ is slightly different from the true \mathbf{F} with diagonal elements of 1.750. We will see that such differences do not affect the convergence of stratified inverses toward the true inverse.

The one-bounce stratified inverse can be easily computed as the inverse of $\hat{\mathbf{F}}$:

$$\mathbf{T}_0^{-1} = \begin{pmatrix} 0.5359 & 0.0000 & 0.0000 & 0.0000 \\ 0.0000 & 0.6120 & 0.0000 & 0.0000 \\ 0.0000 & 0.0000 & 0.5594 & 0.0000 \\ 0.0000 & 0.0000 & 0.0000 & 0.5698 \end{pmatrix}. \quad (37)$$

The two-bounce, to five-bounce stratified inverses are respectively computed as:

$$\mathbf{T}_1^{-1} = \begin{pmatrix} 0.5631 & -0.0070 & -0.0550 & -0.0206 \\ -0.0070 & 0.5606 & -0.0628 & -0.0235 \\ -0.0550 & -0.0628 & 0.5594 & 0.0000 \\ -0.0206 & -0.0235 & 0.0000 & 0.5698 \end{pmatrix}, \quad (38)$$

$$\mathbf{T}_2^{-1} = \begin{pmatrix} 0.5707 & 0.0003 & -0.0571 & -0.0214 \\ 0.0003 & 0.5730 & -0.0568 & -0.0213 \\ -0.0571 & -0.0568 & 0.5715 & 0.0045 \\ -0.0214 & -0.0213 & 0.0045 & 0.5715 \end{pmatrix}, \quad (39)$$

$$\mathbf{T}_3^{-1} = \begin{pmatrix} 0.5713 & -0.0002 & -0.0586 & -0.0220 \\ -0.0002 & 0.5711 & -0.0589 & -0.0220 \\ -0.0586 & -0.0589 & 0.5711 & 0.0044 \\ -0.0220 & -0.0220 & 0.0044 & 0.5714 \end{pmatrix}, \quad (40)$$

and

$$\mathbf{T}_4^{-1} = \begin{pmatrix} 0.5715 & 0.0001 & -0.0586 & -0.0220 \\ 0.0001 & 0.5715 & -0.0586 & -0.0220 \\ -0.0586 & -0.0586 & 0.5715 & 0.0045 \\ -0.0220 & -0.0220 & 0.0045 & 0.5715 \end{pmatrix}. \quad (41)$$

We can see that \mathbf{T}_4^{-1} in (41) is very close to the ground truth inverse in (35). For a diffuse scene, we can compute the first-bounce light transport as

$$\mathbf{F}_{ii} = \frac{1}{(\mathbf{T}_4^{-1})_{ii}}, \quad (42)$$

that is by assigning the diagonal elements of \mathbf{F} as the reciprocal of the corresponding diagonal elements of \mathbf{T}_4^{-1} . This leads to

$$\mathbf{F} = \begin{pmatrix} 1.7498 & 0.0000 & 0.0000 & 0.0000 \\ 0.0000 & 1.7498 & 0.0000 & 0.0000 \\ 0.0000 & 0.0000 & 1.7498 & 0.0000 \\ 0.0000 & 0.0000 & 0.0000 & 1.7498 \end{pmatrix}, \quad (43)$$

which is very close to the true \mathbf{F} . Once we have a good estimate for the true \mathbf{F} (as compared to the initial $\hat{\mathbf{F}}$), we could perform physical light bounce separation as in Seitz et al (2005).

Stratified inverses with an order higher than four can be similarly evaluated. We show the convergence of the stratified inverses against the true inverse in Fig. 5, where the error is measured by the Frobenius norm in log domain $\log \|\mathbf{T}_n^{-1} - \mathbf{T}^{-1}\|_F$. For stratified inverses to have physical interpretation in terms of light bounces, the choice of \mathbf{F} has to be close to the true \mathbf{F} . We can see that our approximated $\hat{\mathbf{F}}$ in (36) is close to the true \mathbf{F} , which is a diagonal matrix with diagonal elements of value 1.75 in this example. Fig. 5 also shows that the convergence curve for using the true \mathbf{F} and our choice of $\hat{\mathbf{F}}$ are very close together. In general, we can also consider \mathbf{F} as a preconditioner in our formulation, where the choice of \mathbf{F} can affect the convergence rate of the stratified inverse. Naive options such as $\mathbf{F} = \mathbf{I}$, which is equivalent to ignoring the preconditioner, are not guaranteed to work, as shown in Fig. 5.

Stratified inverses are only physically meaningful when they are computed with the true \mathbf{F} for a Lambertian scene. If $\hat{\mathbf{F}}$ is not close to the true \mathbf{F} and one wants the full physical interpretation for stratified inverses in terms of light bounces, one could first recover the true \mathbf{F} from stratified inverses using the approximated $\hat{\mathbf{F}}$ and then re-run the iteration with the true \mathbf{F} .

5.2 Computational Time Comparison

To evaluate the computational efficiency of stratified inverses, we generate a 5120×5120 \mathbf{T} matrix with a sparsity of 1.5% non-zero elements from an open Cornell box as shown in Fig. 4(b). Computing the true inverse takes 65 seconds for this \mathbf{T} matrix on Matlab running on a 64-bit machine with 2.67GHz Intel processor and 8 GB RAM. The computational time and accuracy measured with $\log \|\mathbf{T}_i^{-1} - \mathbf{T}^{-1}\|_F$ for the first




Figure 5: Computational performance comparison on a 4×4 \mathbf{T} matrix.




Figure 6: Computational performance comparison on a 5120×5120 \mathbf{T} matrix.

seven stratified inverse terms are shown in Fig. 6 (a). We observe that the stratified inverses converge sharply when we choose \mathbf{F} as the diagonal of \mathbf{T} , while convergence does not happen for $\mathbf{F} = \mathbf{I}$. Furthermore, the computational time for \mathbf{T}_7^{-1} is more than 8 times faster than computing the true inverse.

Fig. 6 (a) indicates the compromise between computational time and accuracy. The balance may depend on the target application, where we may stop the computation when the stratified inverse hits a desired level of accuracy. The accuracy measure $\log \|\mathbf{T}_i^{-1} - \mathbf{T}^{-1}\|_F$ adopted in Fig. 6 (a) is not too useful as we may not know \mathbf{T}^{-1} in practice. The error measure $\|\mathbf{T}\mathbf{T}_n^{-1} - \mathbf{I}\|_F$ offers a better alternative as it can be easily computed given \mathbf{T} . Fig. 6 (b) shows that the error drops as the number of terms in the stratified inverse increases.

We note that very low resolutions were used in order to make comparisons. It was not even possible to run the pseudo-inverse computation (in terms of both computing time and memory) for the larger scenes shown in our results. Also note that there




Figure 7: (a) The experimental setup and the concave wall corner scene. (b) The super-pixel map on the camera image.

is considerable sparsity in the light transport of real scenes, and the stratified inverse computation therefore just involves sparse matrix-matrix multiplication. In these cases, the effective speedup of the stratified inverse method is several orders of magnitude, enabling a full matrix inversion.

6 Results

Having described the theory of the stratified inverse method, we now turn to a demonstration of its practical utility. We focus on radiometric compensation, which is one important application of light transport inversion. For a static scene, we can acquire the transport matrix \mathbf{T} between projector and camera. This tells us how a particular projected image will be affected by interreflections. To compensate the input, and obtain the desired result, we must invert \mathbf{T} , and use this inverse to pre-multiply the projected image. The challenge is inversion of the light transport matrix at reasonable resolutions, and stratified inverses provide a principled way to cancel a certain number of interreflection bounces.

We begin by validating the method on a simple scene with a single wall corner, and then showing results with more complicated geometry and reflectance properties.

6.1 Validation with Simple Wall Corner

The first experimental scene is a concave wall corner as shown in Fig. 7(a), which demonstrates significant light interreflections between the two sides of the wall that join and form a corner. Note that there would be no interreflection if the wall corner was convex. The light interreflections are also evident when an all white image is projected onto the scene as shown in Fig. 8.

For our experimental setup, we used a Canon 450D camera and a Dell 2400MP projector as shown in Fig. 7(a). We linearized the system by first linearizing the camera response using a Macbeth color checker and then linearizing the projector by projecting a sequence of grayscale images with increasing intensity. For simplicity, we consider grayscale light transport where the projector and the camera respectively project and capture grayscale images. As a result, we do not need to model the color mixing matrix




Figure 8: The projector output when projecting a constant-value image and three of its horizontal scan lines. The black-color horizontal is the mean value of projector output.




Figure 9: The accuracy of stratified inverses with respect to the pseudo-inverse in terms of root mean squared error (RMSE)

between the projector and the camera. For accurate scene measurement, we capture high dynamic range images of the scene through a sequence of images with different exposures.

For the first experiment, we focus on a small region centered at the wall corner, so that the corresponding \mathbf{T} is of a manageable size for comparison to explicit pseudo-inverse computation. We group 4×4 projector pixels into a super-pixel and restrict the active region in the projector to be of dimensions 31×51 super-pixels at the center. The corresponding super-pixel on the camera is a group of camera pixels that a projector super-pixel has the maximum response on. The intensity of a camera super-pixel corresponds to the mean intensity of the camera pixels in the group. The super-pixel map on the camera for the wall corner scene is shown in Fig. 7(b). With this setting, \mathbf{T} is a square matrix of dimension 1938×1938 .

We also perform a real radiometric compensation by projecting a non-negative im-

age obtained by

$$\mathbf{I}_{\text{in}} = f_+(\mathbf{T}^{-1}\mathbf{I}_{\text{desired}}), \quad (44)$$

where the function $f_+(\cdot) = \max(\cdot, 0)$ truncates values to 0 since the projector cannot display negative values, and $\mathbf{I}_{\text{desired}}$ is the image that we desire to see in the scene.

6.1.1 Data Acquisition

In this simple experiment, we choose the brute-force method for light transport acquisition, as it gives us an accurate \mathbf{T} . After having acquired \mathbf{T} , we evaluate its quality by comparing the image obtained by projecting a uniform intensity image $\mathbf{I}_{\text{const}}$ with that generated by simulation using the acquired \mathbf{T} , i.e., by comparing with $\mathbf{I}_{\text{out}} = \mathbf{T}\mathbf{I}_{\text{const}}$. We plot three horizontal scan lines from the respective images as shown in Fig. 8. We can see that the scan lines coincide nicely in shape, although the simulated scan lines look noisier, which could be due to the hard grouping of super-pixels. However, the similarity of the scan line graphs in shape indicates that \mathbf{T} is a sufficiently good model of the scene. Note that the scan lines peak at the center region, which is due to the significant light interreflections between the sides of the wall that form an L shape. Overall, the scan lines are far from being constant intensity, although a constant intensity image is projected.

6.1.2 Stratified Inverse

For stratified inverse computation, we first extract the first-bounce light transport matrix \mathbf{F} from \mathbf{T} as the diagonal of \mathbf{F} . Knowing both \mathbf{T} and \mathbf{F} , we can compute stratified inverses using (14). We compute the stratified inverses of \mathbf{T} up to four light bounces (or four terms in the series) as well as the pseudo-inverse of \mathbf{T} . In this case, we consider the pseudo-inverse as the ground-truth and compute the error of these stratified inverses with respect to the ground truth in terms of root mean squared error (RMSE).

Fig. 9 shows how the error of stratified inverses decreases as the number of light bounces or terms increases. We observe that the error becomes sufficiently low from two light bounces onward.

6.1.3 Radiometric Compensation

Next, we evaluate the capability of stratified inverses for projector radiometric compensation. In this experiment, we specify the desired image $\mathbf{I}_{\text{desired}}$ to be a uniform intensity image. Radiometric compensation is expected to flatten the fluctuation shown in Fig. 8 as much as possible, as it undoes the light interreflections in the scene. We can easily simulate the effect of radiometric compensation as below

$$\mathbf{I}_{\text{out}} = \mathbf{T}f_+(\mathbf{T}^{-1}\mathbf{I}_{\text{desired}}) \simeq \mathbf{I}_{\text{desired}}. \quad (45)$$

The function $f_+(\cdot) = \max(\cdot, 0)$ truncates values to 0, since the projector cannot display negative values. It has little effect in practice, but is necessary to respect physical constraints.

The results are shown in Fig. 10(a). Note that the intensity scale on the y-axis has been expanded in contrast to that in Fig. 8. We note that the scan line compensated




Figure 10: Results for the case of a constant-value $\mathbf{I}_{\text{desired}}$. (a) Horizontal scan lines for the simulated \mathbf{I}_{out} . (b) Horizontal scan lines for the captured \mathbf{I}_{out} . For both (a) and (b), the number i on the scan lines indicates i -bounce stratified inverse. The black-color scan line is one for the true inverse.

using the 4-bounce stratified inverse is quite close to the desired constant intensity scan line. Visually, we also see that the scan lines compensated using stratified inverses get closer to the desired constant intensity scan line as the number of light bounces goes up.

Fig. 11 shows the inputs to the projector computed using 1-bounce and 4-bounce stratified inverses, and the pseudo-inverse, as well as the corresponding outputs observed by the camera. Note the projector inputs for the 4-bounce stratified inverse and pseudo-inverse are similar. The black dots in the projector inputs are due to the clipping function f_+ that ensures the nonnegativity of the projector inputs. Also note that scene interreflections are still visible in the case of the 1-bounce stratified inverse, while the outputs for the 4-bounce stratified inverse and pseudo-inverse are almost of uniform intensity and close to the desired output.

We show the corresponding central horizontal scan lines of the output image seen by the camera in Fig. 10(b). Comparing to the simulated results in Fig. 10(a), the intensity fluctuation in Fig. 10(b) is slightly larger. We can measure the fluctuation with the ratio of the scanline disparity (i.e., $\max - \min$) to the desired constant value. For the 4-bounce scan line in Fig. 10(b), the fluctuation measure is about 0.3, which is significantly smaller than that of the uncompensated scan lines in Fig. 8 which is around 0.9. Similar to the simulation, we also see that the scan lines in Fig. 10(b) get closer to the desired constant intensity scan line as the number of light bounces increases.

6.2 Complex Scenes

To evaluate our method, we perform an experiment on a complex scene of a surface tiled with an array of hemispheres as shown in Fig. 1(a). The scene interreflections are significant at the concave regions between the hemispheres, and the output image before compensation is seriously distorted both geometrically and radiometrically. For this complex scene, we acquire a light transport matrix of slightly higher resolution using a brute-force method with the projector having 92×116 super-pixels of dimension




Figure 11: Results for the experiment where the desired image is a uniform intensity image. (Top row) The inputs to the projector for various kinds of matrix inversion. (Middle row) The outputs observed by the camera for various kinds of matrix inversion. We extract a scan line from each of the outputs and show them in Fig. 10. (Bottom row) The desired constant intensity image.

4×4 pixels. In this case, \mathbf{T} has a dimension of 10672×10672 . Fig. 1(e) shows the complex desired image displayed in high resolution. Note that when there is no radiometric compensation on the projector output, as shown in Fig. 1(b), interreflections are visible between the hemispheres which results in the reduced contrast, and there is geometry distortion in the image. After radiometric compensation with the 1-bounce stratified inverse, the geometry distortion is in general removed as shown in Fig. 1(c) (some slight local distortion in the image is due the coarse super-pixel structure). We can see that the scene interreflections are significantly reduced after the 4-bounce compensation as shown in Fig. 1(d). The compensation also rectifies the distortion observed in the uncompensated image.

In terms of processing time required, computing the 4-bounce stratified inverses requires about 8 seconds, while computing pseudo-inverse becomes impractical. We believe the few seconds of compute time to invert the full \mathbf{T} matrix is easily justifiable for a static scene, especially considering the time for acquisition of the transport matrix.




Figure 12: *The high-resolution results for the wall corner scene. Subfigures from (a) to (d) show increasing fidelity. The uncompensated projector output in (a) shows geometric distortion and scene interreflections. The projector output (b) by the 1-bounce stratified inverse is compensated mainly in geometry. The projector output (c) is compensated both geometrically and radiometrically by the 4-bounce stratified inverse. The result (c) is close to the desired image (d). Note that the minor intensity differences between the desired image in (d) and the projector outputs (c) are due to the residual photometric nonlinearity left uncorrected by the simplistic projector response curve.*

New images can now be compensated with a simple matrix-vector multiply.

6.3 High-resolution Compensation

To assess our method for a high-resolution case, we acquire a high-resolution light transport matrix with the projector having 384×512 super-pixels of dimension 2×2 pixels for the convex wall corner scene. For efficiency and ease of implementation, \mathbf{T} is acquired using the stripe-scanning method in Ding et al (2009), instead of the brute-force method. In this case, \mathbf{T} is a matrix with dimension of 196608×196608 . We specify a complex city-scene desired image as shown in Fig. 12(a). Before radiometric compensation, the output image appears to be distorted and has strong scene interreflections as can be seen in Fig. 12(b). Note that in the region with strong scene interreflections the intensity contrast is much reduced. The compensation from both the 1-bounce and the 4-bounce stratified inverses rectifies the image distortion as shown in Fig. 12(b) and (c) respectively, and the scene interreflections are more much subdued for the 4-bounce solution. Also note that the compensated result in Fig. 12(c) closely matches the desired output in Fig. 12(d). Note that computing the 4-bounce stratified inverses requires about 40 seconds, while computing the pseudo-inverse is impractical at these resolutions.

6.4 Limitations of Our Method

As we can see in Fig. 10, the compensated scan lines oscillate about the desired scan line before settling down. The intensity oscillation represents both over-compensation and under-compensation. We find that such oscillation is slow to settle down at the sharp transitions or folds in physical scenes. For example, the wall corner scene of Fig. 12 has a visible crack at the wall corner. Intensity oscillation remains at the fold even though the 4-bounce compensation has compensated most parts of the scene. Whereas in the scene of Fig. 1, there are shadow regions around some of the hemi-

spheres, which is a form of sharp scene transition. Similarly, the attempt in compensating for the shadow regions results in white ring artifacts which is a form of overcompensation even though other parts of the scene have been compensated. In future work, we will look into ways to speed up the convergence for stratified compensation at sharp scene transitions.

7 Conclusions and Future Work

We proposed a stratified computational framework for computing the inverse light transport. More importantly, we showed that the stratified inverse corresponds to canceling the corresponding light bounce for each term. As opposed to our previous work in Ng et al (2009) that was built on the interreflections cancellation operator introduced in Seitz et al (2005), we relate stratified inverses directly to Kajiya’s rendering equation (Kajiya, 1986). We validate our theoretical results on simulations that show the accuracy and computational efficiency of stratified inverses. We also show the application of stratified inverses for projector radiometric compensation.

There are many important directions of future work. In terms of efficiency, the use of hierarchical or wavelet matrix representations can enhance the sparsity and speed up higher bounces of the matrix multiplications in our methods. In terms of acquisition, we would like to extend the work of Nayar et al (2006) to full light transport separation, even in the presence of high-frequency local effects like sharp subsurface interactions. Using low-frequency illumination to acquire the global light transport promises to speed up acquisition, and point the way towards radiometric compensation for dynamic scenes.

We have presented only one application in projector radiometric compensation, but \mathbf{T} and its inverse are crucial in many other problems in computer graphics and vision—for all of which our theoretical development should provide new insights. For example, inverse light transport can also be used for light bounce separation (Seitz et al, 2005) and shape estimation (Liu et al, 2010). Given the initial work (Liu et al, 2010) showing the connection between inverse light transport and inverse rendering (Marschner, 1998; Ramamoorthi and Hanrahan, 2001), there may be further applications for inverse light transport. Moreover, large-scale matrix inversion is also crucial in other problem domains such as Google’s PageRank vector computation. In future work, we expect many other application domains to benefit from the stratified inverse framework.

8 Acknowledgements

We are grateful to Tony Quee-Seng Quek for giving feedbacks on an earlier draft of the manuscript. We wish to thank Manmohan Chandraker and Mao-Pei Tsui for their great suggestions on various technical issues. We would also like to thank Yasuyuki Matsushita, Jiaping Wang and Yue Dong of Microsoft Research Asia for their discussion and suggestions on experimental techniques. Finally, we would like to thank Zhiyong Huang and Joo Hwee Lim for their support on this project. Ramanpreet Singh Pahwa and Jiamin Bai are supported by National Science Scholarships awarded by A*STAR

Graduate Academy of Singapore. Ravi Ramamoorthi is supported by ONR YIP grant N00014-10-1-0032, as well as NSF CAREER Award 09-24968 and ONR PECASE grant N00014-09-1-0741.

References

Arvo J, Torrance K, Smits B (1994) A framework for the analysis of error in global illumination algorithms. In: Proceedings of ACM SIGGRAPH, pp 75–84

Ashdown M, Okabe T, Sato I, Sato Y (2006) Robust content-dependent photometric projector compensation. In: Proceedings of IEEE International Workshop on Projector Camera Systems

Bai J, Chandraker M, Ng TT, Ramamoorthi R (2010) A dual theory of inverse and forward light transport. In: Proceedings of European Conference on Computer Vision

Bimber O (2006) Emerging Technologies of Augmented Reality: Interfaces and Design, Idea Group Pub., chap Projector-based augmentation, pp 64–89

Bimber O, Emmerling A, Klemmer T (2005) Embedded entertainment with smart projectors. IEEE Computer Magazine 38(1):48–55

Bimber O, Grundhofer A, Zeidler T, Danch D, Kapakos P (2006) Compensating indirect scattering for immersive and semi-immersive projection displays. In: Proceedings of the IEEE Virtual Reality Conference, pp 151–158

Debevec P, Hawkins T, Tchou C, Duiker H, Sarokin W, Sagar M (2000) Acquiring the reflectance field of a human face. In: Proceedings of ACM SIGGRAPH, pp 145–156

Ding Y, Xiao J, Tan KH, Yu J (2009) Catadioptric projectors. In: Proceedings of Computer Vision and Pattern Recognition

Fujii K, Grossberg M, Nayar S (2005) A projector-camera system with real-time photometric adaptation for dynamic environments. In: Proceedings of Computer Vision and Pattern Recognition

Gortler S, Schröder P, Cohen M, Hanrahan P (1993) Wavelet radiosity. In: Proceedings of ACM SIGGRAPH, pp 221–230

Habe H, Saeki N, Matsuyama T (2007) Inter-reflection compensation for immersive projection display. In: Proceedings of Computer Vision and Pattern Recognition

Hanrahan P, Salzman D, Aupperle L (1991) A rapid hierarchical radiosity algorithm. In: Proceedings of ACM SIGGRAPH, pp 197–206

Kajiya JT (1986) The rendering equation. In: Proceedings of ACM SIGGRAPH, pp 143–150

Liu S, Ng TT, Matsushita Y (2010) Shape from second-bounce of light transport. In: Proceedings of European Conference on Computer Vision

Majumder A, Gopi M, Seales BW, Fuchs H (1999) Geometric stitching for real-time panoramic image generation using texture maps. In: Proceedings of ACM International Conference on Multimedia

Marschner SR (1998) Inverse rendering for computer graphics. PhD thesis, PhD Thesis, Cornell University

Mukaigawa Y, Kakinuma T, Ohta Y (2006) Analytical compensation of inter-reflection for pattern projection. In: Proceedings of the ACM symposium on Virtual reality software and technology, p 268

Nayar S, Peri H, Grossberg M, Belhumeur P (2003) A projection system with radiometric compensation for screen imperfections. In: Proceedings of IEEE International Workshop on Projector Camera Systems

Nayar S, Krishnan G, Grossberg M, Raskar R (2006) Fast separation of direct and global components of a scene using high frequency illumination. ACM Transactions on Graphics 25(3):935–944

Ng R, Ramamoorthi R, Hanrahan P (2003) All-frequency shadows using non-linear wavelet lighting approximation. Proceedings of ACM SIGGRAPH 22(3):376–381

Ng TT, Pahwa RS, Bai J, Quek QS, , Tan KH (2009) Radiometric compensation using stratified inverses. In: Proceedings of International Conference on Computer Vision

O'Toole M, Kutulakos KN (2010) Optical computing for fast light transport analysis. In: Proceedings of ACM SIGGRAPH Asia

Peers P, Mahajan D, Lamond B, Ghosh A, Matusik W, Ramamoorthi R, Debevec P (2009) Compressive light transport sensing. ACM Transactions on Graphics 28(1)

Ramamoorthi R (2007) Precomputation-based rendering. Foundations and Trends® in Computer Graphics and Vision 3(4):281–369

Ramamoorthi R, Hanrahan P (2001) A signal-processing framework for inverse rendering. In: Proceedings of ACM SIGGRAPH, pp 117–128

Raskar R, Welch G, Cutts M, Lake A, Stesin L, Fuchs H (1998) The office of the future: A unified approach to image-based modeling and spatially immersive displays. In: Proceedings of ACM SIGGRAPH, pp 179–188

Raskar R, Welch G, Fuchs H (1999) Spatially augmented reality. AK Peters, Ltd.

Seitz S, Matsushita Y, Kutulakos K (2005) A theory of inverse light transport. In: Proceedings of International Conference on Computer Vision, pp 1440–1447

Sen P, Darabi S (2009) Compressive dual photography. Computer Graphics Forum 28(2):609–618

Sen P, Chen B, Garg G, Marschner S, Horowitz M, Levoy M, Lensch H (2005) Dual photography. ACM Transactions on Graphics 24(3):745–755

Song P, Cham T (2005) A theory for photometric self-calibration of multiple overlapping projectors and cameras. In: Proceedings of IEEE International Workshop on Projector Camera Systems

Wang J, Dong Y, Tong X, Lin Z, Guo B (2009) Kernel nyström method for light transport. ACM Transactions on Graphics 28(3)

Wetzstein G, Bimber O (2007) Radiometric compensation through inverse light transport. In: Proceedings of Pacific conference on computer graphics and applications, pp 391–399

Yang R, Majumder A, Brown M (2004) Camera based calibration techniques for seamless flexible multi-projector displays. In: Proceedings of Applications of Computer Vision Workshop

Yu Y, Debevec P, Malik J, Hawkins T (1999) Inverse global illumination: Recovering reflectance models of real scenes from photographs. In: Proceedings of ACM SIGGRAPH, pp 215–224